Washington University in St. Louis

Engineering

Network Technology Adoption Why It Is Hard (From Concepts to Realities)

Roch Guérin

Department of Computer Science & Engineering

Washington University in Saint Louis

December 2018

Acknowledgments

- This talk is based on discussions and collaborations with many colleagues, post-docs, and students over the years. In particular,
  - » Hadi Afrasiabi, Wells Fargo Securities
  - » Constantine Dovrolis, Georgia Tech.
  - » Kartik Hosanagar, Wharton Business School, University of Pennsylvania
  - » Youngmi Jin, KDDI
  - » Mehdi Nikkhah, 6D.ai
  - » Andrew Odlyzo, School of Mathematics, University of Minnesota
  - » Soumya Sen, Carlson School of Management, University of Minnesota
  - » Zhi-Li Zhang, Department of Computer Science, University of Minnesota

Any errors and omissions in this talk are, however, all mine

Where It All Started

- Before moving to academia, I spent about 12 years at the IBM T.J. Watson Research Center working on a wide range of networking technologies (and attending the IETF and the ATM Forum)
  - » While I learned a lot and we developed some pretty nifty systems, most did not really gain widespread adoption
- This triggered some questions to try to understand why
  - » The history of networking is filled with technologically sound defunct companies (Bay Networks, Cabletron, Cascade, FORE, Ipsilon, Newbridge,…) and standards (ATM, IPX, Token Ring, XTP,…)
- The “simple” answer it is that technical quality is a necessary but far from sufficient condition for success
  - » Many factors and complex interactions can and will influence the outcome
- This talk tries to shed some light on these complex phenomena

Outline

- A quick overview of basic adoption models and some of the strange outcomes that can arise
  - » Utility functions, adoption decisions and dynamics, and equilibria
- The case of IPv6
  - » Stake-holders, dependencies, empirical studies, and an attempt at reverse engineering how things evolved
- IETF standards (RFCs)
  - » RFCs are akin to breadcrumbs that track the Internet's evolution
  - » Our goal: Identifying key features present across them and apply statistical analysis to isolate factors likely to play an important role in a protocol's success or failure
- A few brief concluding remarks in an attempt to summarize lessons learned

Modeling (Network) Technology Adoption

- Basic framework assumes some form of rationality on the part of users/adopters
  - » We adopt only if there is a benefit to us, and if several options are feasible we pick the one with the highest benefit
- Utility function $U_i(t)$ measures the benefit derived by user i when adopting technology T at time t
  - » $U_i(t) = f_i$(what I like in T at time t) – g_i(what I don't like in T at time t)
  - » User i adopts iff $U_i(t) > 0$
    - Both $f_i(.)$ and $g_i(.)$ can vary with i and t, as well as a function of how many other users adopt T (positive and negative *externalities*)
- **Externalities are a trade-mark of networking technologies and one of the reasons their adoption is hard to predict**

A "Basic" Example

Positive externality — Negative externality — Intrinsic value — Cost

$$U_i(x(t)) = \alpha_i x - \frac{\beta_i x(t)}{1 - x(t)} + v - c(x(t))$$

- $U_i(x(t))$: Utility of user i when $x(t)$ users have adopted
- Linear positive externality (Metcalfe's law) – $\alpha_i x(t)$
- Delay-like negative externality – $\beta_i x(t)/(1 - x(t))$
- Fixed intrinsic value v and fixed or log-like cost $c(x(t))$

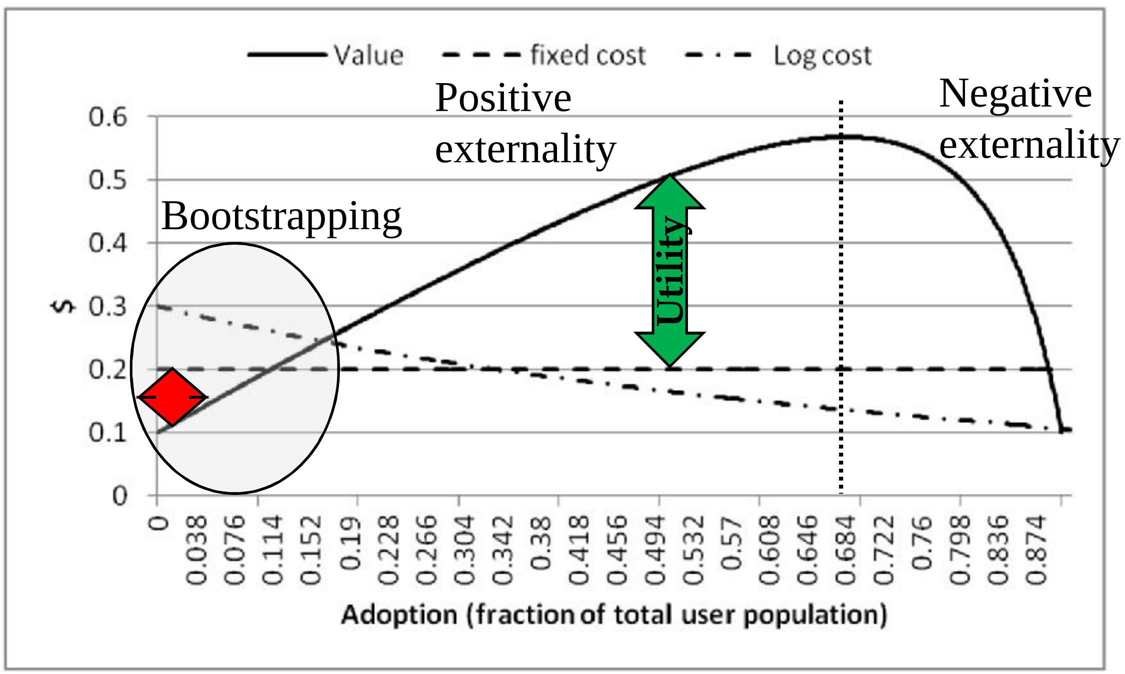

- Assumes homogeneous users for simplicity

Adoption Dynamics – Overview

- Given an adopter's population, how does adoption ***evolve***?
  - » Can we ***influence*** it, *e.g.*, through pricing or incentives
- Adoption at time *t*: $x(t)$
  - » Given $x(t)$, $H(x(t))$ is the number of users who should adopt
    - At equilibrium $H(x^*) = x^*$
- Adoption dynamics
  - » A diffusion process with a *rate* $\gamma < 1$

$$\frac{dx(t)}{dt} = \gamma\left(H\big(x(t)\big) - x(t)\right)$$

- Identifying equilibria and adoption trajectories then boils down to solving systems of differential equations

In General: Standard Adoption Patterns (for successful technologies)

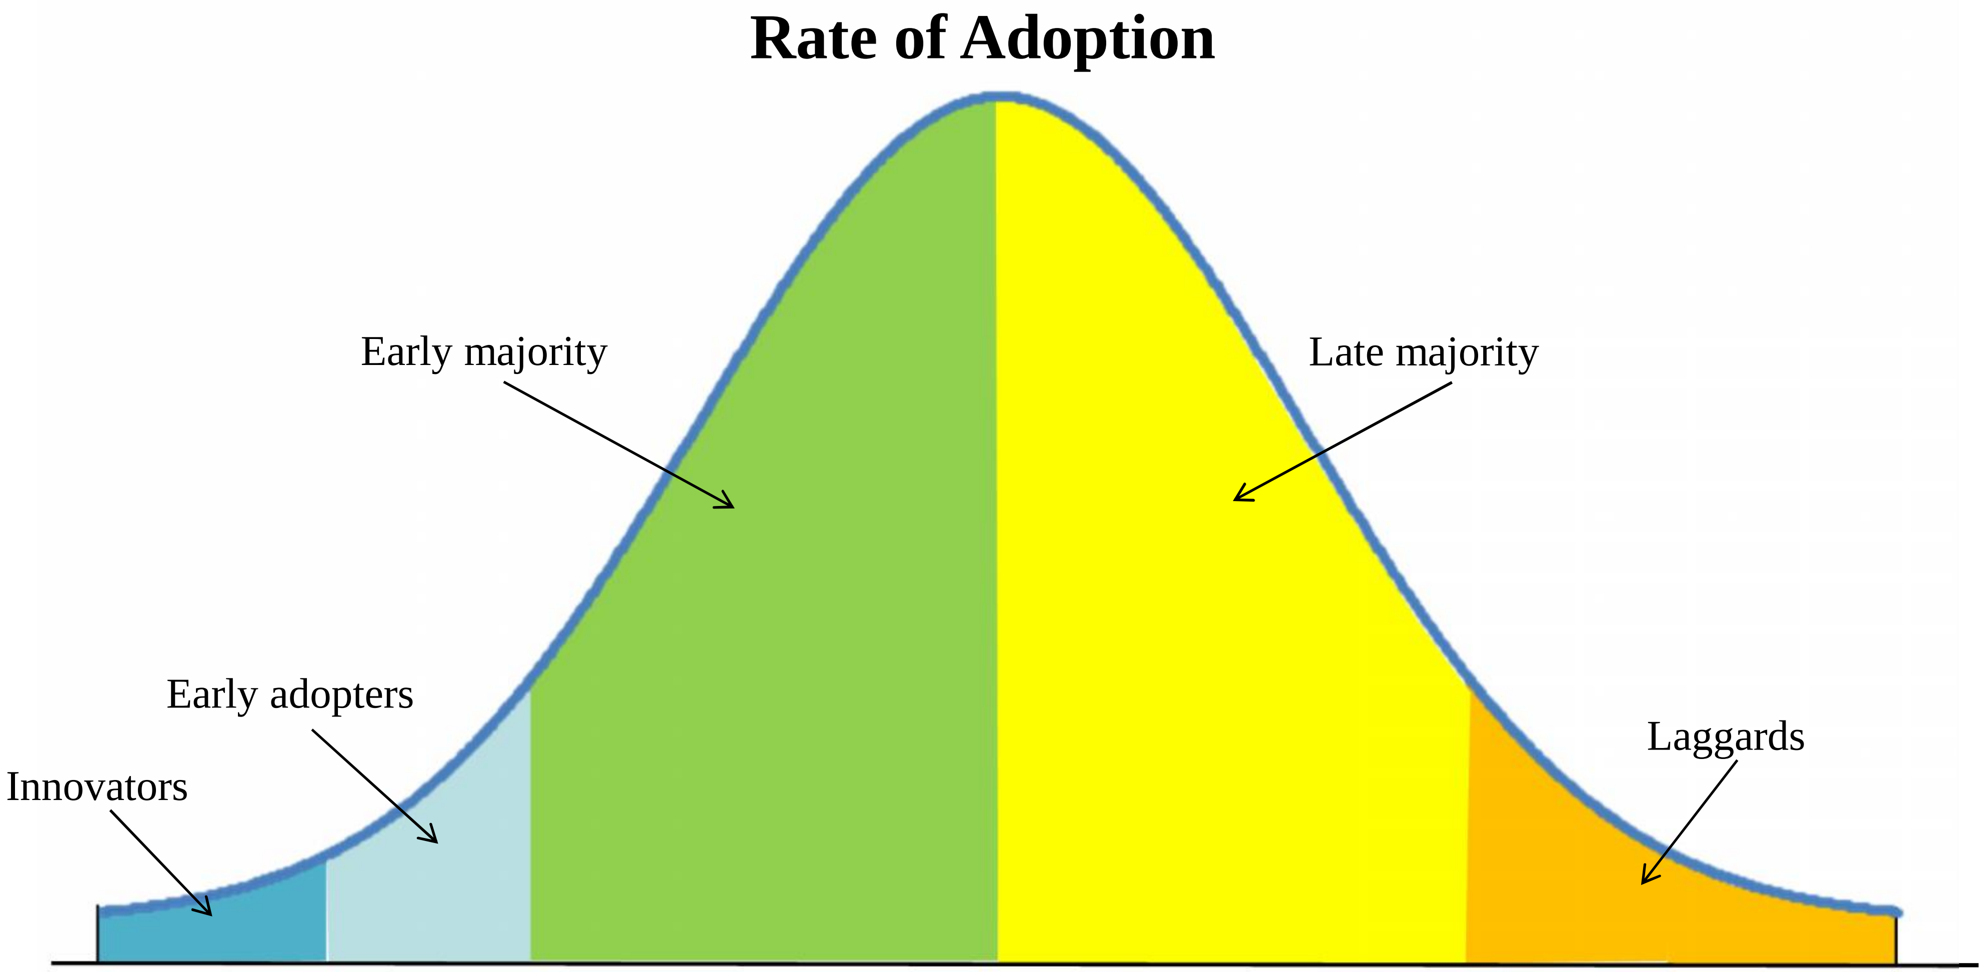

A Common Networking Scenario

- Two competing technologies with gateways for inter-operability
- Outcome depends on
  - » Initial adoption level of the incumbent
    - New technology “wins” only if incumbent is not too entrenched
  - » Quality of gateways
    - They can help or hurt the new technology depending on system parameters
- Most importantly, in practice the outcome is hard to predict from observations

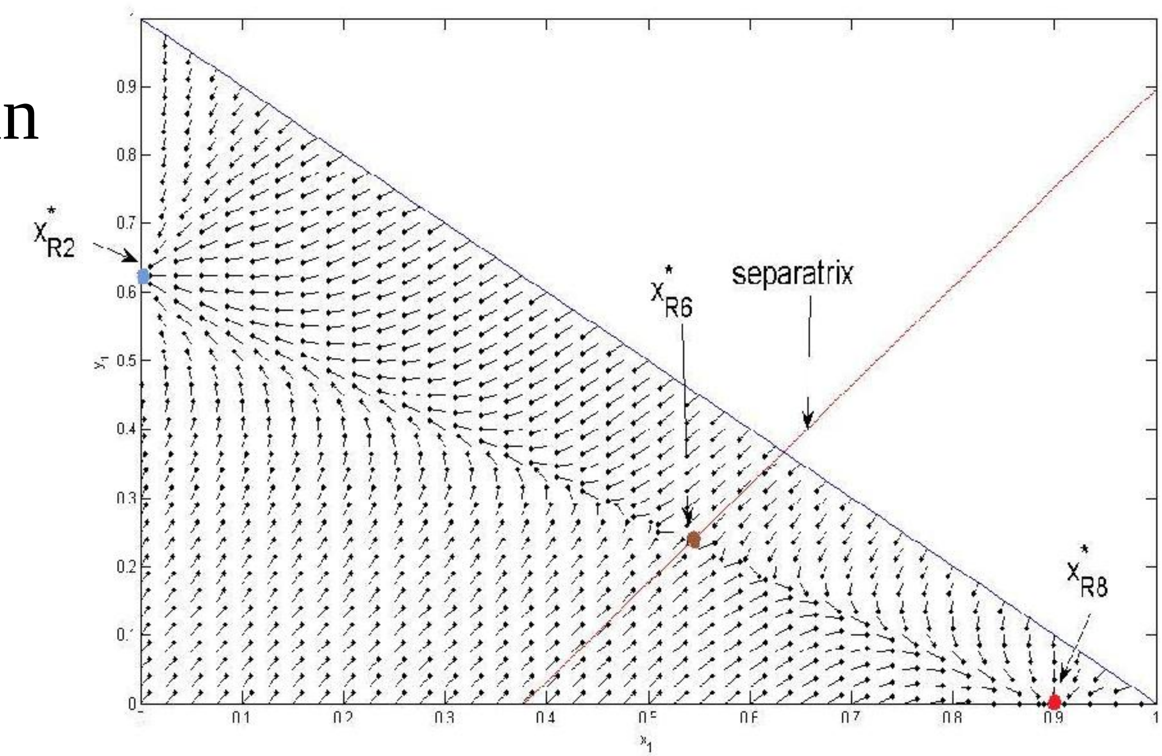

In Summary

- We have tools at our disposal to explore adoption scenarios
- Network technologies are particularly challenging because of
  - Externalities, gateways, etc.

  That all contribute non-linearities that make predicting outcomes difficult
- Models can help forecast obstacles and predict trends, but can only offer ***broad guidelines <u>not</u> design recipes***

Outline

- A quick overview of basic adoption models and some of the strange outcomes that can arise
  - Utility functions, adoption decisions and dynamics, and equilibria
- The case of IPv6
  - Stake-holders, dependencies, empirical studies, and an attempt at reverse engineering how things evolved
- IETF standards (RFCs)
  - RFCs are akin to breadcrumbs that track the Internet's evolution
  - Our goal: Identifying key features present across them and apply statistical analysis to isolate factors likely to play an important role in a protocol's success or failure
- A few brief concluding remarks in an attempt to summarize lessons learned

The Case of IPv6

- IPv6 adoption is now "under way", BUT it took much longer than anticipated (it was standardized over 20 years ago)
- **Why has it been so hard,** and could we have made it better?

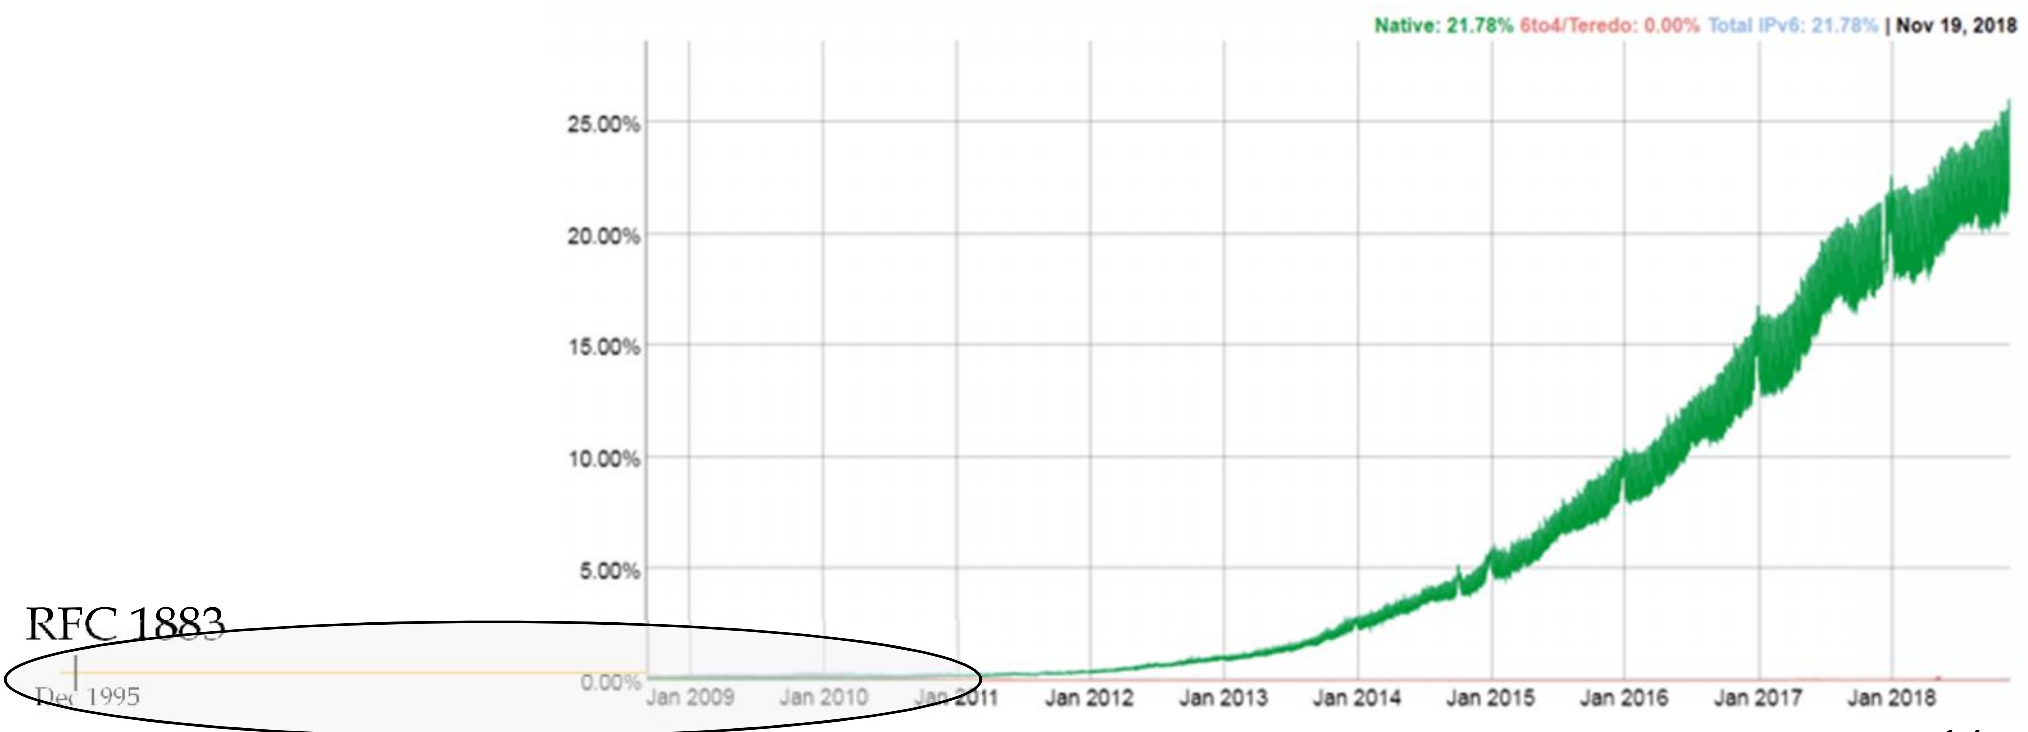

https://www.google.com/intl/en/ipv6/statistics.html

Exploring This Question

1. Reviewing the ecosystem
   » Stakeholders: Internet users, technology vendors, content providers, service providers
   » Decision factors and their role: Cost, technology and connectivity (IPv6) quality, dependencies on other stake-holders, etc.
2. Gathering data
   » From external sources (CAIDA, Google, company websites, etc.)
   » From home-grown monitoring project
      • Tracking top 1M+ websites for IPv6 connectivity and quality differential
3. Interpreting data
   » Cause and effect relationships
4. A simple model for "validation" purposes

The Internet Ecosystem – Passive & Active Players

1. **Internet Users** ✗
   » Mostly oblivious to technology, *i.e.*, whether it's IPv4 or IPv6, but affected by availability of applications and content, as well as connectivity quality
2. **Internet technology vendors** ✓
   » Concerned about market growth and development costs
3. **Internet Content Providers (ICPs)** ✓
   » Focus on delivering content (and ads) to Internet users
4. **Internet Service Providers (ISPs)** ✓
   » Deliver Internet connectivity and grow user-base, but concerned about cost (both capital and operational) and quality

- Internet users play little direct role but can influence IPv6 adoption decisions by other stakeholders (externalities)

IPv6 Adoption – Decision Factors

- IPv6 value is its larger address space and not much else… So,

- Internet Technology Developers (ITDs) adoption
  - » Driven by demand (from ICPs & ISPs)
    - Demand must offset development costs
- Internet Content Providers (ICPs) adoption
  - » Driven by (enough) IPv6 eyeballs and benefits/quality of IPv6 connectivity
    - **A strong externality factor** (grows with # of IPv6 users and IPv6 quality)
- Internet Service Providers (ISPs) adoption
  - » Driven by (IPv4) address acquisition costs and migration costs (operational and translation IPv6 $\leftrightarrow$ IPv4)
    - **A strong externality factor**: # IPv6 users and ICPs

Data (Ours & Others) to Track IPv6 Evolution

e.g., see https://www.internetsociety.org/resources/2018/state-of-ipv6-deployment-2018/

- ITDs: Routers, OSes, and applications
  - » IPv6 readiness and performance (compared to IPv4)
- ICPs: Content from top websites (Alexa ranking)
  - » IPv6 websites accessibility and benefits (compared to IPv4)
- ISPs: IPv6 footprint over the public Internet
  - » Number of ASes advertising IPv6 prefixes, number of IPv6 peering links, number of IPv6 routes
- What does the data tell us?

High-Level Findings

A three phase evolution across stakeholders

- **Phase 1 [1995-2009]** (Stagnation): Marginal availability and/or immature technology
- **Phase 2 [2009-2011]** (Emergence): Telltale signs of early adoption and greater maturity
- **Phase 3 [2011-]** (Acceleration): Still not mainstream, but a growing tangible footprint
  - » Or to use our earlier adoption terminology, we are now in the "**early majority**" stage*

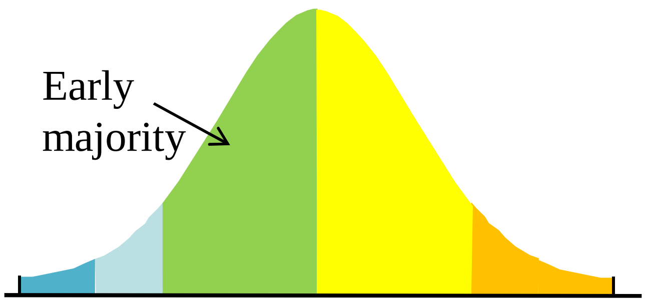

* See also https://www.internetsociety.org/resources/2018/state-of-ipv6-deployment-2018/

Representative Data – ITDs

- Evolution of IPv6 availability & quality
  - » Pre 2009: Dismal IPv6 forwarding performance, few if any IPv6 capable applications, OS support rife with glitches
  - » After 2009: IPv6 forwarding problems mostly solved, most OSes with mature IPv6 support (2011)
  - » Today: Availability and performance problems mostly a thing of the past
    - Application support, though, is not yet ubiquitous (a long tail)

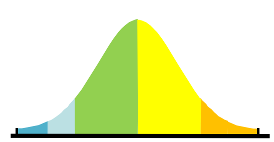

From https://technet.microsoft.com/en-us/library/hh994905(d=printer).aspx

Internet Services

Service name	Full IPv6 support	Limited IPv6 Support	Full IPv4 support	Additional information
Azure	●		●	Azure Networking announcements for Ignite 2016 IPv6 for Azure VMs available in most regions
Bing			●	
Dynamics CRM Online			●	
Microsoft.com	●		●	
Office 365	●		●	IPv6 support in Office 365 services Office 365 IPv6 Test Plan for Chiayi Educational Network Center
OneDrive			●	
Outlook.com			●	
Skype		● Mobile Only		
SQL Azure			●	
System Center Advisor			●	
Microsoft Intune			●	
Windows Update	●		●	
Xbox.com	●		●	
Yammer			●	

Representative Data – ISPs

- IPv6 (#ASes, peering links, routes) evolution
  - » Marginal prior to 2009
  - » Telltale signs of growth between 2009 and 2011
  - » A noticeable pick-up of pace since 2011
- A similar picture when it comes to traffic and DNS queries

From CAIDA http://goo.gl/OhqWNM

	IPv4 ASes	IPv6 ASes	IPv4 Peering	IPv6 Peering
2009	23k	515	50k	1904
2011	29k	1183	78k	2738
2013	34k	2419	109k	8881

From Geoff Huston http://bgp.potaroo.net/v6/as6447/

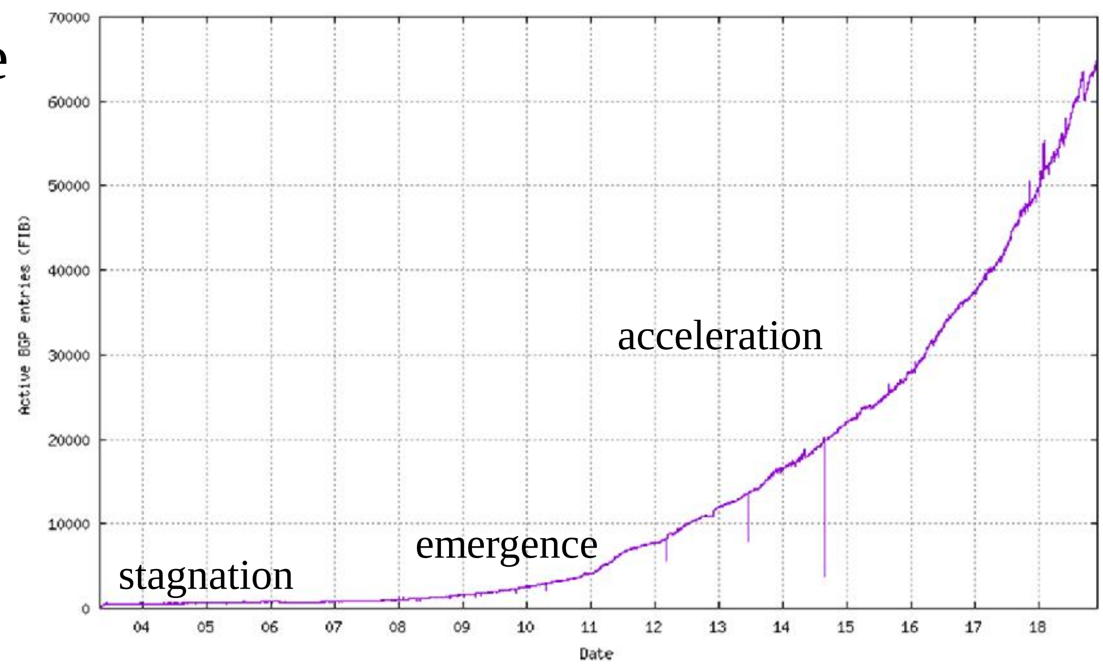

Representative Data – ICPs

- ICPs IPv6 accessibility
  - » Little to no adoption before 2009 among top 1M websites
  - » Adoption started taking off after 2009, especially among more popular sites
    - Top 1k sites at ~14% in 2014 vs. 28% five years later
    - Top 1M now at 17%*
  - » But progress towards 100% is likely to take time

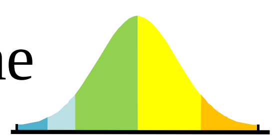

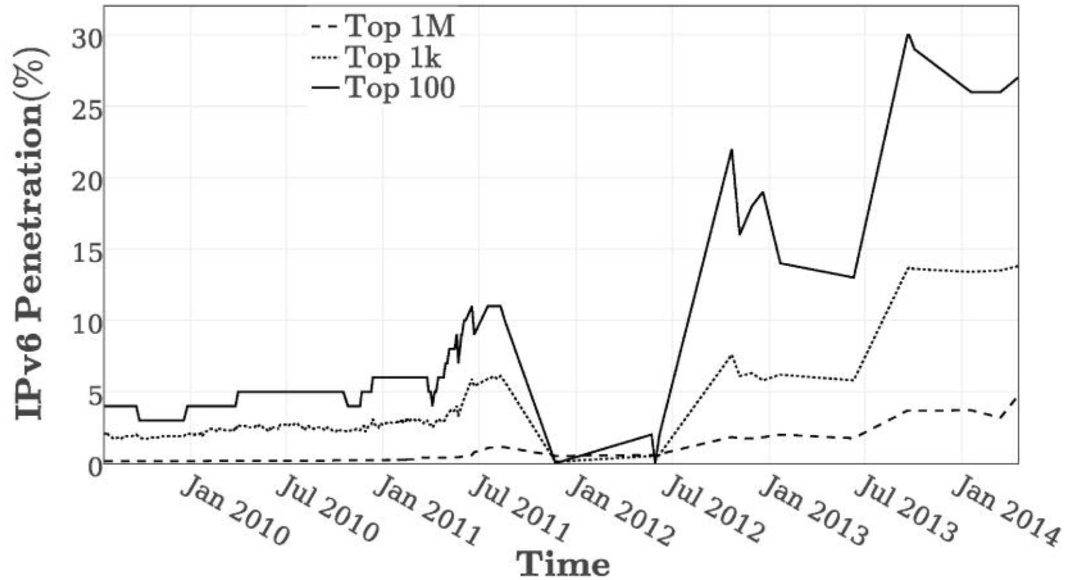

From https://www.worldipv6launch.org/measurements/

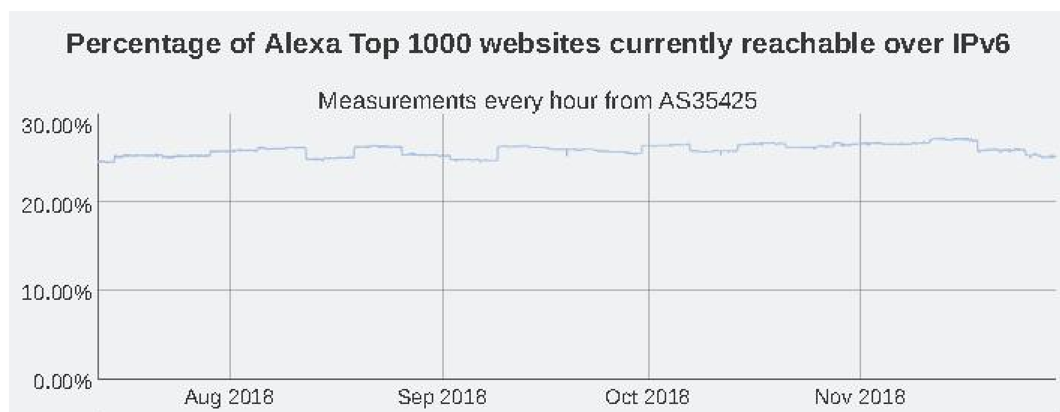

* https://www.internetsociety.org/resources/2018/state-of-ipv6-deployment-2018/

An Attempt at Explaining Observations

- Influential factors

Factors (increases)	Impact on Utility		
	ISPs	**ICPs**	**ITDs**
IPv6 demand	x	x	+
IPv4 address cost	–	x	x
IPv6 upgrade cost	–	–	x
Translation cost	~	x	x
Size of IPv6 user base	~	+	~
Number of IPv6 ICPs	+	~	~
IPv6 quality	x	+	x

x : No impact
+ : Positive impact
– : Negative impact
~ : Marginal impact

- What changed over time?
  - » IPv6 **demand**
    - We finally ran out of IPv4 addresses
    - Growing awareness through government mandates and events like IPv6 Launch Day & World IPv6 Day
    - IPv4 addresses now cost real money
  - » IPv6 **quality**
    - Technology parity with IPv4 not until 2009
    - End-to-end parity took longer (2011)

IPv4 Addresses Are Not Free Anymore!

http://ipv4marketgroup.com/ipv4-pricing/
July 18, 2018

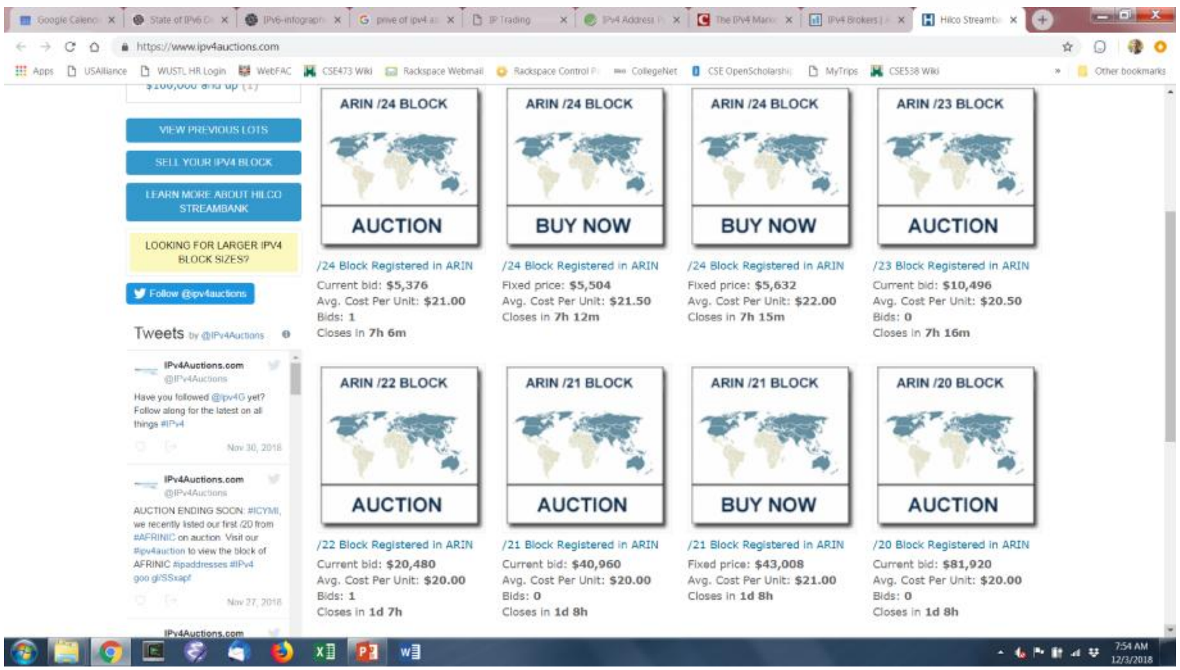

https://www.ipv4auctions.com/
December 3, 2018

An Attempt at Explaining Observations

- Influential factors

Factors (increases)	Impact on Utility		
	ISPs	**ICPs**	**ITDs**
IPv6 demand	x	x	+
IPv4 address cost	–	x	x
IPv6 upgrade cost	–	–	x
Translation cost	~	x	x
Size of IPv6 user base	~	+	~
Number of IPv6 ICPs	+	~	~
IPv6 quality	x	+	x

x : No impact
+ : Positive impact
– : Negative impact
~ : Marginal impact

- What changed over time?
  - » IPv6 **demand**
    - We finally ran out of IPv4 addresses
    - Growing awareness through government mandates and events like IPv6 Launch Day & World IPv6 Day
    - IPv4 addresses now cost real money
  - » IPv6 **quality**
    - Technology parity with IPv4 not until 2009
    - End-to-end parity took longer (2011)

IPv6 and IPv4 Performance Parity

- In 2009 barely 40% of websites could be reached (from the US) over IPv6 with quality on par with that of an IPv4 connection
- This improved to 80% by 2013, with the remaining 20% a toss-up
- Two primary causes of performance differences
  1. Packet forwarding performance
  2. Paths to destination
  - » The first was fixed circa 2009, while eliminating or shortening IPv6 "detours" took longer

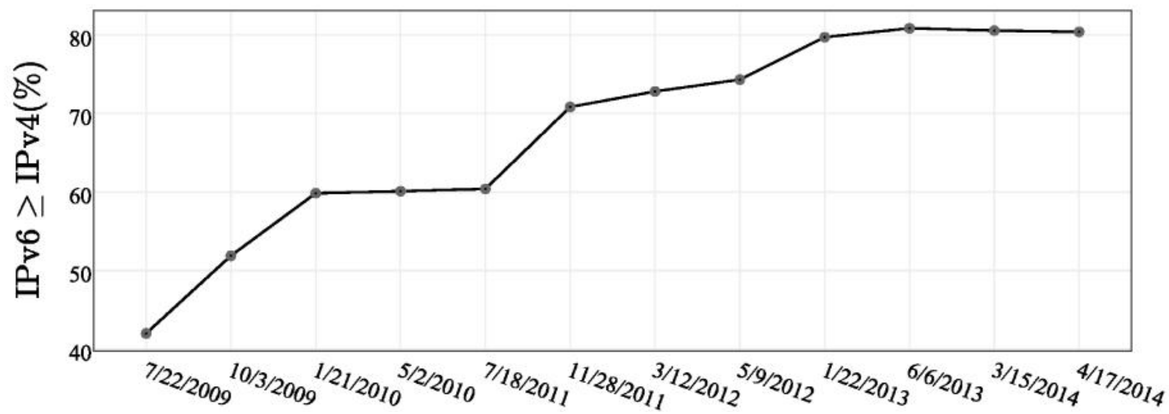

% of websites with IPv6 connectivity on par or better than IPv4

IPv6 ≥ IPv4	Top 100k sites		Top 1M sites	
	2011	2013	2011	2013
Same path	94%	100%	90%	94%
Different paths	70%	79%	74%	84%

Summarizing Causes and Effects

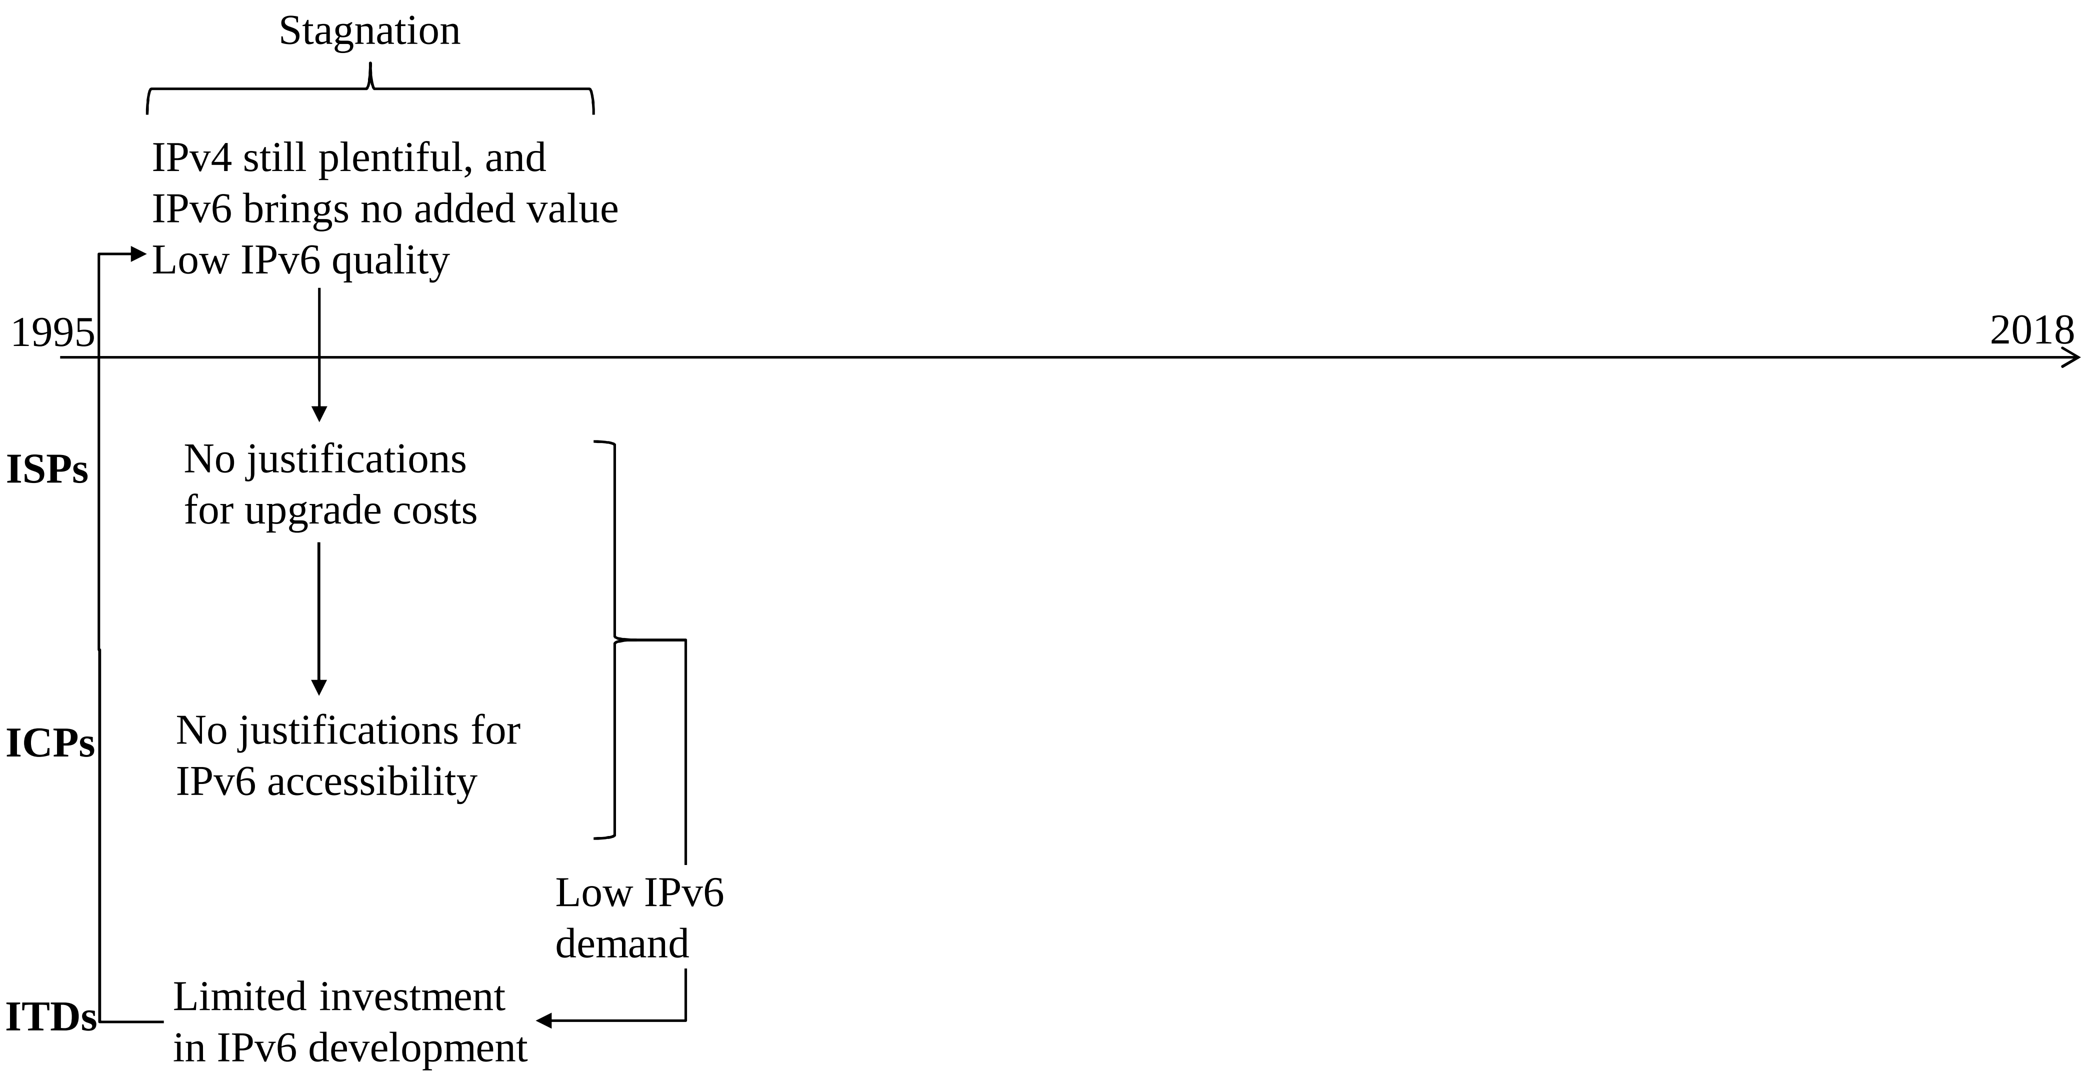

Summarizing Causes and Effects

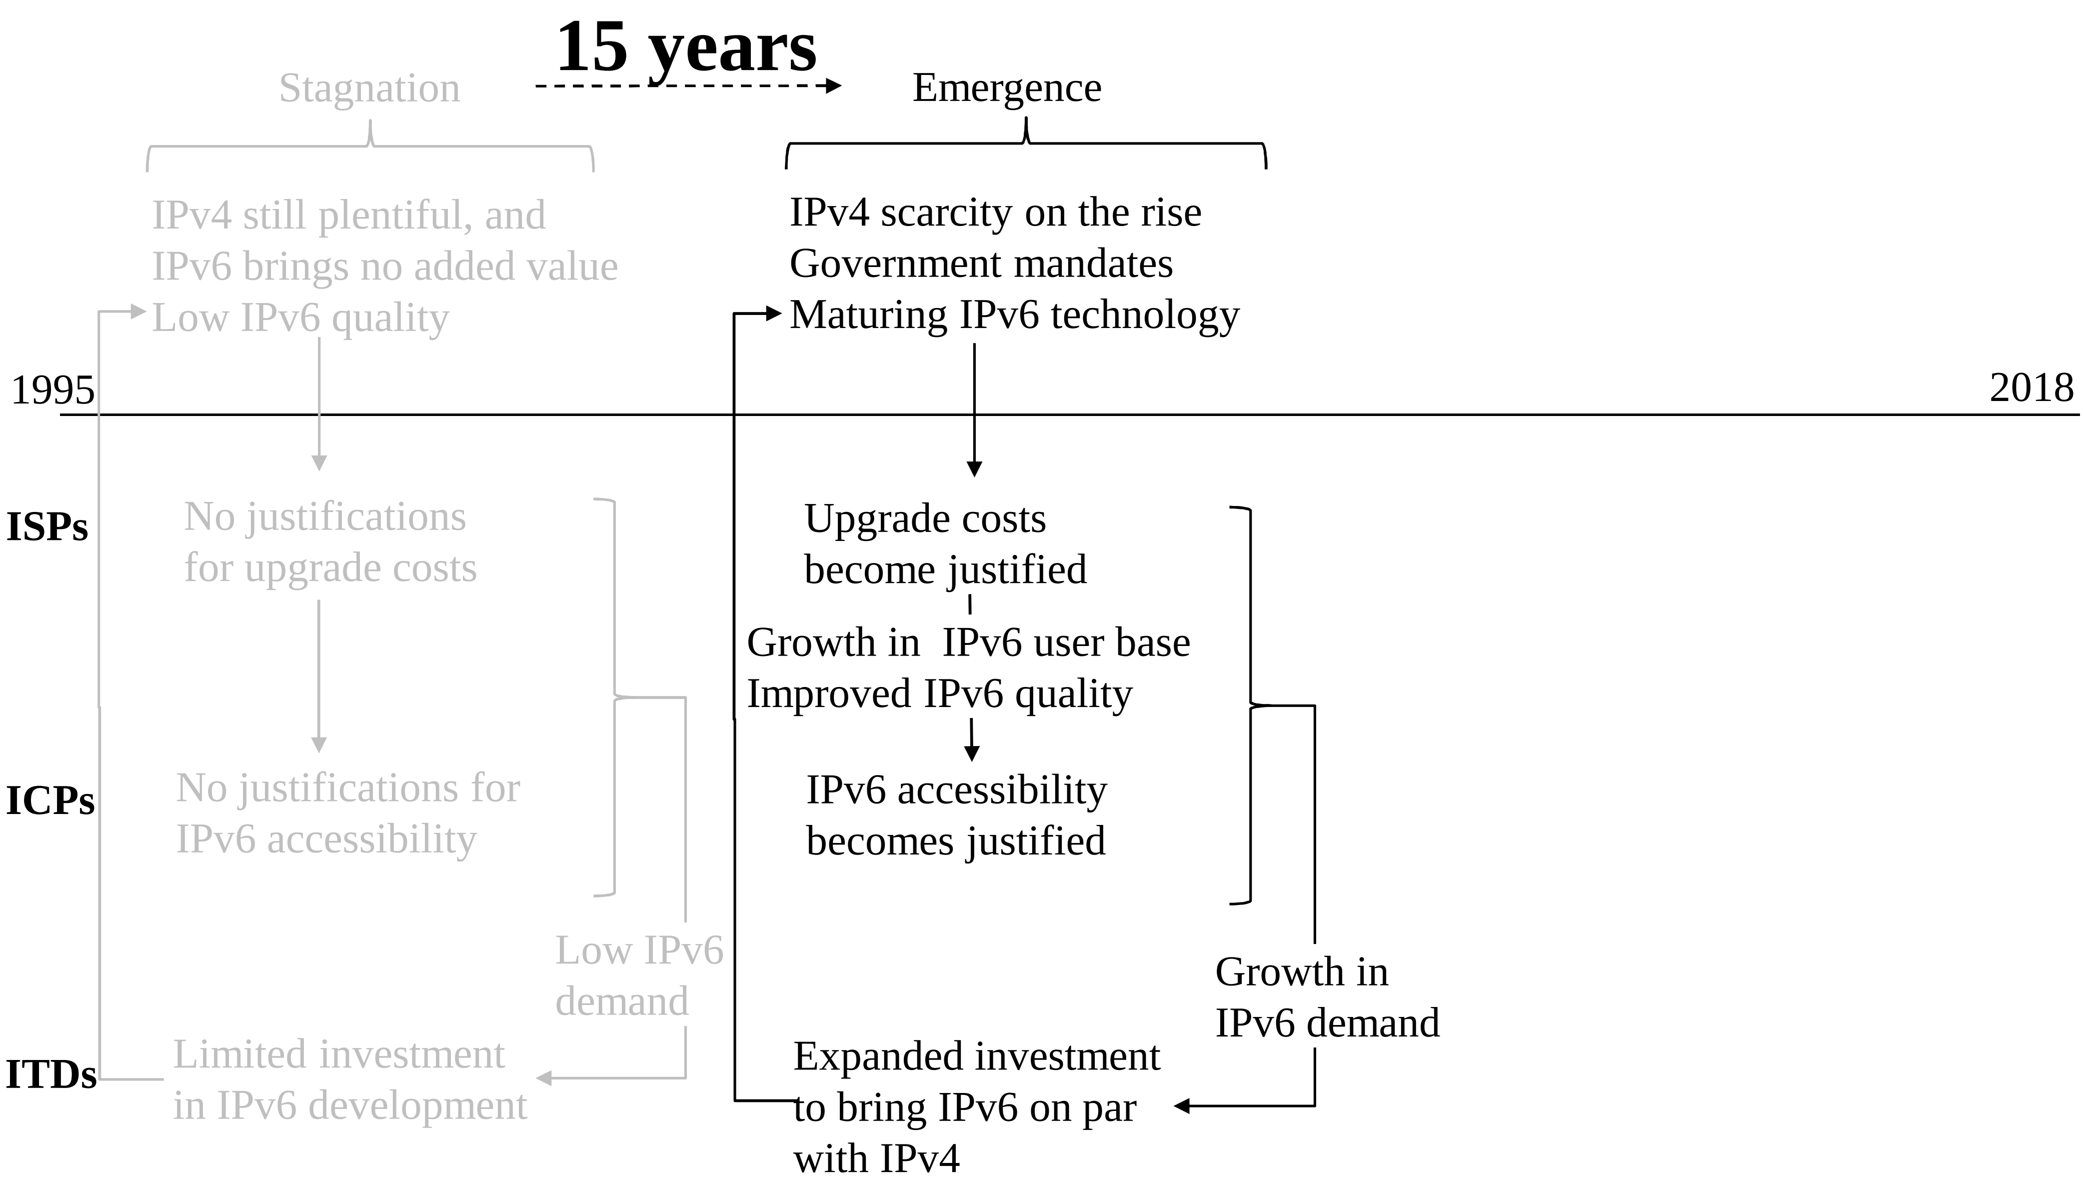

Summarizing Causes and Effects

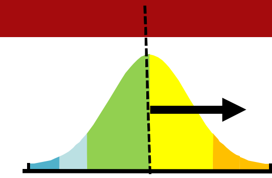

20+ years to full adoption

1995 — 2018

Stagnation (15 years →)
- IPv4 still plentiful, and IPv6 brings no added value
- Low IPv6 quality

ISPs: No justifications for upgrade costs
ICPs: No justifications for IPv6 accessibility
Low IPv6 demand
ITDs: Limited investment in IPv6 development

Emergence (2-5 years → Impact of externalities)
- IPv4 scarcity on the rise
- Government mandates
- Maturing IPv6 technology

ISPs: Upgrade costs become justified
Growth in IPv6 user base
Improved IPv6 quality
ICPs: IPv6 accessibility becomes justified
Growth in IPv6 demand
ITDs: Expanded investment to bring IPv6 on par with IPv4

Acceleration (→)
- IPv4 exhaustion
- IPv6 on par with IPv4

ISPs: IPv6 upgrade a necessity
Continued growth in IPv6 user base
IPv6 on par with IPv4
ICPs: IPv6 accessibility becomes necessary
Steady growth in IPv6 demand
ITDs: IPv6 investment now steady-state

What Could We Have Done Better/Differently?

- Three basic options
  1. Increase incentives for early IPv6 adoption
  2. Decrease disincentives for early IPv6 adoption
  3. Increase disincentives for IPv4

  (1–2: **Positive actions**; 3: **Negative actions**)
- Option 1
  - » Make IPv6 more "valuable" than IPv4, *e.g.*, better support for security or mobility
    - But it's unclear how much sway this would have had in the early days of the Internet, where connectivity was the primary value it offered
- Option 2
  - » Make IPv6 backward compatible with IPv4 to facilitate migration
    - Lower the cost of upgrading/migrating to IPv6
  - » Ensure parity with IPv4 (in terms of quality and stability) on day 1
    - Incentives to ITDs to invest more in IPv6 development (forwarding, routing, DNS, OS support, etc.)
- Option 3
  - » Make IPv4 more expensive or lower quality than IPv6 (add penalty to IPv4)
    - Charge for IPv4 address, but make IPv6 addresses free
    - Give IPv6 traffic precedence over IPv4
  - » But this could have been at the cost of jeopardizing the Internet success (there were competitors)

Outline

- A quick overview of basic adoption models and some of the strange outcomes that can arise
  - » Utility functions, adoption decisions and dynamics, and equilibria
- The case of IPv6
  - » Stake-holders, dependencies, empirical studies, and an attempt at reverse engineering how things evolved
- IETF standards (RFCs)
  - » RFCs are akin to breadcrumbs that track the Internet’s evolution
  - » Our goal: Identifying key features present across them and apply statistical analysis to isolate factors likely to play an important role in a protocol’s success or failure
- A few brief concluding remarks in an attempt to summarize lessons learned

Towards a Better Understanding of Protocol Adoption in the Internet

- Models can pinpoint key factors, but as IPv6 illustrated, the impact of small quantitative changes are hard to predict

- But the Internet is now ~50 years old (RFC 1 was published in 1969), and we have many protocol examples to study from

⇒ A data driven approach

» Targeted at Internet Standards and Proposed Standards
  - About 3300 Standards Track RFCs

» Can we infer likely success or failure factors, and their relative impact?
  - A random sample of 200 RFCs yields a success rate of about 60%

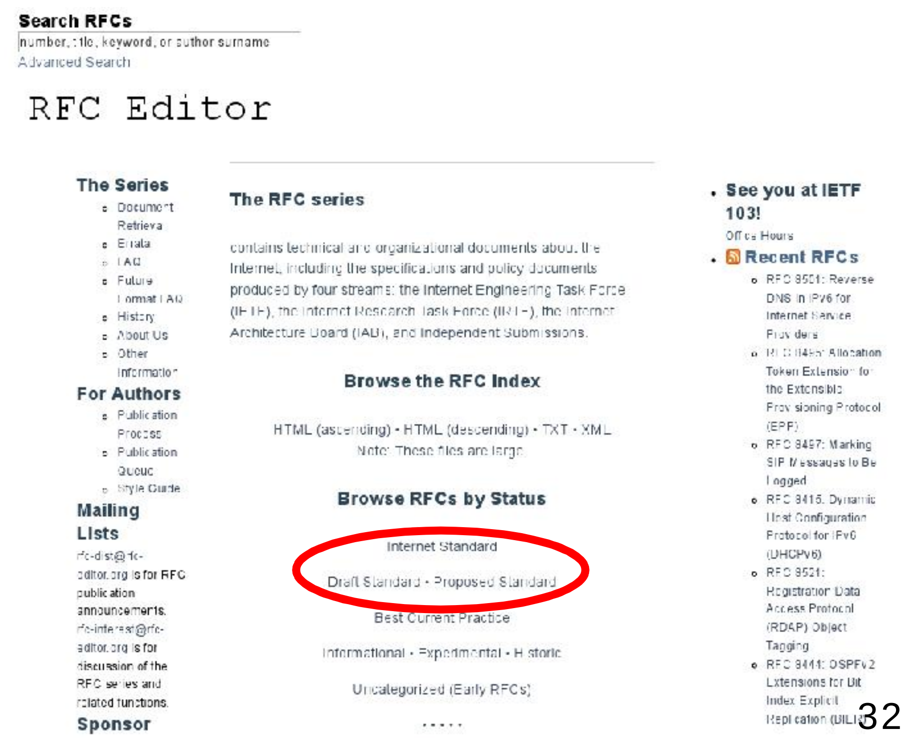

Data Collection Overview

- Start with ~2500 RFCs that have become Internet Standards or Proposed Standards
  - » Distributed across 6 major areas:
    - Application & real-time (ART, ~30%)
    - Operations & management (OPS, ~15%)
    - Internet (INT, ~20%)
    - Transport (TSV, ~8%)
    - Routing (RTG, ~14%)
    - Security (SEC, ~13%)
- Sample each area to preserve distribution of RFCs areas
  - » An observational case control study: Positive (successful RFCs), negative (failed RFCs)
  - » Initial data set of about 450 RFCs (about 251 after "cleanup")
- Classify each RFC according to binary or categorical features, including "area," success or failure, and various other factors
  - » Success ≡ widespread adoption among target audience
  - » Classification categories and results available at http://goo.gl/r3KP2K

	RFC Title	RFC Date	RFC IETF Status	Widely Adopted? [1]	IETF Area [2]	RFC Type [3]	Incumbent? [4]	Scope of Adoption [5]	Value Increase? [6]	Change Others? [7]	Add Value to Upstream? [8]	Security? [9]	Scalability? [10]	Performance? [11]
RFC 854	Telnet Protocol Specification	5/1/1983	Internet Standard	Yes [12]	art	N	No	E2E	Yes	No	No	No	No	No
RFC 1184	Telnet Linemode Option	10/1/1990	Proposed Standard	Yes [13]	art	EB	-	E2E	Yes	No	No	No	No	No
RFC 1647	TN3270 Enhancements	7/1/1994	Proposed Standard	Yes [14]	art	EB	- [15]	E2E	Yes	No	No	No	No	No
RFC 1730	Internet Message Access Protocol - Version 4	12/1/1994	Proposed Standard	Yes [16]	art	EB	- [17]	E2E	No	No	Yes [18]	No	No	No
RFC 1731	IMAP4 Authentication Mechanisms	12/1/1994	Proposed Standard	Yes [19]	art	EB	-	BN	No	No	No	Yes	No	No
RFC 1782	TFTP Option Extension	3/1/1995	Proposed Standard	Yes [20]	art	EB	-	E2E	Yes	No	Yes [21]	No	No	No
RFC 1869	SMTP Service Extensions	11/1/1995	Proposed Standard	Yes [22]	art	EB	-	E2E	Yes [23]	No	No	No	No	No
RFC 1939	Post Office Protocol - Version 3	5/1/1996	Proposed Standard	Yes [24]	art	N [25]	No	E2E	No	No	No	No	No	No
RFC 2045	Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies	11/1/1996	Draft Standard	Yes [26]	art	N	No	E2E	Yes	No	Yes [27]	No	No	No
RFC 2068	Hypertext Transfer Protocol -- HTTP/1.1	1/1/1997	Proposed Standard	Yes [28]	art	EB	- [29]	E2E	Yes	No	No	Yes	No	Yes
RFC 2163	Using the Internet DNS to Distribute MIXER Conformant Global Address Mapping (MCGAM)	1/1/1998	Proposed Standard	No [30]	art	N	No	UBN	Yes	No	No [31]	No	No	No
RFC 2244	ACAP -- Application Configuration Access Protocol	11/1/1997	Proposed Standard	No [32]	art	NI	Yes [33]	E2E	Yes	No	No	No	No	No
RFC 2371	Transaction Internet Protocol Version 3.0	7/1/1998	Proposed Standard	No [34]	art	N [35]	No	E2E	Yes	No	No	No	No	No
RFC 2389	Feature negotiation mechanism for the File Transfer Protocol	8/1/1998	Proposed Standard	Yes [36]	art	EB	-	E2E	Yes	No	No	No	No	No
RFC 2447	iCalendar Message-Based Interoperability Protocol (iMIP)	11/1/1998	Proposed Standard	Yes [37]	art	N	No	E2E	Yes [38]	No	No	No	No	No
RFC 2651	The Architecture of the Common Indexing Protocol (CIP)	8/1/1999	Proposed Standard	No [39]	art	NI	Yes [40]	BN	Yes [41]	No	No	No	No	No
RFC 2793	RTP Payload for Text Conversation	5/1/2000	Proposed Standard	Yes [42]	art	EB	-	E2E	Yes	No	No	No	No	No
RFC 2821	Simple Mail Transfer Protocol	4/1/2001	Proposed Standard	Yes [43]	art	N	No	E2E	Yes	No	No	No	No	No
RFC 2911	Internet Printing Protocol/1.1: Model and Semantics	9/1/2000	Proposed Standard	Yes [44]	art	EB	- [45]	BN	Yes [46]	No	No	No	No	No
RFC 2915	The Naming Authority Pointer (NAPTR) DNS Resource Record	9/1/2000	Proposed Standard	Yes [47]	art	EB	-	UBN	Yes	No	Yes [48]	No	No	No
RFC 2920 a.k.a. STD 60	SMTP Service Extension for Command Pipelining	9/1/2000	Internet Standard	Yes [49]	art	EB	-	E2E	Yes	No	No	No	No	Yes [50]
RFC 2950	Telnet Encryption: CAST-128 64 bit Cipher Feedback	9/1/2000	Proposed Standard	No [51]	art	EB	-	E2E	Yes	No	No	Yes [52]	No	No
RFC 3015	Megaco Protocol Version 1.0	11/1/2000	Proposed Standard	Yes [53]	art	EB	- [54]	BN	No	No	Yes [55]	No	Yes [56]	No
RFC 3080	The Blocks Extensible Exchange Protocol Core	3/1/2001	Proposed Standard	No [57]	art	N	No	E2E	Yes	No	No	No	No	No
RFC 3207	SMTP Service Extension for Secure SMTP over Transport Layer Security	2/1/2002	Proposed Standard	Yes [58]	art	EB	-	E2E	Yes	No	No	Yes	No	No
RFC 3261	SIP: Session Initiation Protocol	6/1/2002	Proposed Standard	Yes [59]	art	N	No	E2E	Yes	No	Yes [60]	No	No	No
RFC 3340	The Application Exchange Core	7/1/2002	Historic (changed from Proposed Standard)	No [61]	art	NI	Yes [62]	E2E	Yes	No	No	No	No	No
RFC 3403	Dynamic Delegation Discovery System (DDDS) Part Three: The Domain Name System (DNS) Database	10/1/2002	Proposed Standard	No [63]	art	N	No	UBN	Yes	No	No	No	No	No
RFC 3550	RTP: A Transport Protocol for Real-Time Applications	7/1/2003	Proposed Standard	Yes [64]	art	N	No	E2E	Yes [65]	No	Yes [66]	No	No	No
RFC 4918	HTTP Extensions for Web Distributed Authoring and Versioning (WebDAV)	6/1/2007	Proposed Standard	Yes [67]	art	EB	-	E2E	Yes	No	No	No	No	No
RFC 3659	Extensions to FTP	3/1/2007	Proposed Standard	Yes [68]	art	EB	-	E2E	Yes	No	No	No	No	No
RFC 3730	Extensible Provisioning Protocol (EPP)	3/1/2004	Proposed Standard	Yes [69]	art	N	No	E2E	Yes [70]	No	No	No	No	No
RFC 3885	SMTP Service Extension for Message Tracking	9/1/2004	Proposed Standard	Yes [71]	art	EB	-	E2E	Yes	No	No	No	No	No
RFC 3887	Message Tracking Query Protocol	9/1/2004	Proposed Standard	No [72]	art	N	No	E2E	Yes	No	Yes [73]	No	No	No
RFC 3911	The Session Initiation Protocol (SIP) "Join" Header	10/1/2004	Proposed Standard	No [74]	art	EB	- [75]	E2E	Yes	No	No	No	No	No
RFC 3920	Extensible Messaging and Presence Protocol (XMPP): Core	10/1/2004	Proposed Standard	Yes [76]	art	N	No	E2E	Yes	No	No	No	No	No
RFC 3965	A Simple Mode of Facsimile Using Internet Mail	12/1/2004	Draft Standard	Yes [77]	art	N	No	E2E	Yes	No	Yes [78]	No	No	No
RFC 3977	Network News Transfer Protocol (NNTP)	10/1/2006	Proposed Standard	Yes [79]	art	NI	Yes [80]	E2E	Yes	No	No	No	No	Yes [81]
RFC 3981	IRIS: The Internet Registry Information Service (IRIS) Core Protocol	1/1/2005	Proposed Standard	No [82]	art	NI	Yes	UBN	Yes	No	No	No	No	No

http://goo.gl/r3KP2K

Classification factors

	RFC Title	RFC Date	RFC IETF Status	Widely Adopted? [1]	IETF Area [2]	RFC Type [3]	Incumbent? [4]	Scope of Adoption [5]	Value Increase? [6]	Change Others? [7]	Add Value to Upstream? [8]	Security? [9]	Scalability? [10]	Performance? [11]
RFC 854	Telnet Protocol Specification	5/1/1983	Internet Standard	Yes [12]	art	N	No	E2E	Yes	No	No	No	No	No
RFC 1184	Telnet Linemode Option	10/1/1990	Proposed Standard	Yes [13]	art	EB	-	E2E	Yes	No	No	No	No	No
RFC 1647	TN3270 Enhancements	7/1/1994	Proposed Standard	Yes [14]	art	EB	- [15]	E2E	Yes	No	No	No	No	No
RFC 1730	Internet Message Access Protocol - Version 4	12/1/1994	Proposed Standard	Yes [16]	art	EB	- [17]	E2E	No	No	Yes [18]	No	No	No
RFC 1731	IMAP4 Authentication Mechanisms	12/1/1994	Proposed Standard	Yes [19]	art	EB	-	BN	No	No	No	Yes	No	No

http://goo.gl/r3KP2K

Methodology

- Logistic regression
  - » Because of our small dataset and goal of identifying “risk factors”
  - » It also offers prediction odds instead of binary answer
- Two different tools: SAS and R (consistent results with both)
  - » Forward-backward stepwise regression (p-value threshold of 0.1) and stepwise regressions with Akaike Information Criterion (AIC) to avoid over-fitting and remove insignificant factors
  - » Accuracy evaluated using leave-one-out (LOO) cross-validation and synthetic data (to assess predictive ability on small dataset)
  - » Robustness analysis by adding “noise” to classification process
- Analysis applied to both full dataset and per-area dataset

Findings – Full Dataset

- Strong contributor to likelihood of ***success***
  - » **Adding value to upper layer protocols**
- Strong contributors to likelihood of ***failure***
  - » **Facing an incumbent**
  - » **Lack of backward compatibility with earlier version**
- Nothing overly surprising and mostly consistent with intuition

Findings – Per Area Datasets

- Three pairs of areas were found to be statistically similar
  - » Applications & Security (ART & SEC)
  - » Internet & Transport (INT & TSV)
  - » Operations and Routing (OPS & RTG)

  These groupings are again intuitive
- Key ***success*** and ***failure*** factors

±	ART & SEC	INT & TSV	OPS & RTG
⊕	Adding value to upper layer protocols	Adding value to upper layer protocols	Improving security and scalability
⊖	Non-backward compatible extension or new protocol with incumbent Internet-wide adoption	Non-backward compatible extension or new protocol with incumbent Internet-wide adoption Changes to other protocols	NA

What Is This Telling Us?

- Not all protocols affected by the same factors
- **ART**, **SEC**, **INT**, and **TSV** protocols
  - » Success positively affected by contributions to making other protocols valuable
    - Do they benefit the Internet as a whole?
  - » Main challenges come from **externalities**
- **OPS** & **RTG** protocols
  - » Focus on Internet infrastructure makes impact of security and scalability intuitive
    - Little to no sensitivity to other factors (never faced?)

±	ART & SEC	INT & TSV	OPS & RTG
⊕	Adding value to upper layer protocols	Adding value to upper layer protocols	Improving security and scalability
⊖	Non-backward compatible extension or new protocol with incumbent Internet-wide adoption	Non-backward compatible extension or new protocol with incumbent Internet-wide adoption Changes to other protocols	NA

How Does This Map to IPv6?

- IPv6 part of the Internet (INT) area
  - » No real added value beyond IPv4
  - » Not backward compatible (with IPv4)
  - » Calls for Internet-wide adoption
  - » Requires changes to existing protocols

±	**INT & TSV**
⊕	Adding value to upper layer protocols
⊖	Non-backward compatible extension or new protocol with incumbent Internet-wide adoption Changes to other protocols

- From our findings, IPv6 clearly faced an uphill challenge (our model gave it a **1.5% success probability**…)
- Experimenting with possible alternatives ("what-if" scenarios, ignoring technical feasibility and cost)
  - » Biggest impact is removing "changes to other protocols"
    - Unlikely to be feasible given address expansion from 32 bits to 128 bits
  - » Backward compatibility and incremental deployment also helped
    - A protocol with all three constraints removed had a **75% chance of success**

Outline

- A quick overview of basic adoption models and some of the strange outcomes that can arise
  - » Utility functions, adoption decisions and dynamics, and equilibria
- The case of IPv6
  - » Stake-holders, dependencies, empirical studies, and an attempt at reverse engineering how things evolved
- IETF standards (RFCs)
  - » RFCs are akin to breadcrumbs that track the Internet’s evolution
  - » Our goal: Identifying key features present across them and apply statistical analysis to isolate factors likely to play an important role in a protocol’s success or failure
- A few brief concluding remarks to summarize lessons learned

Conclusions

- Predicting the success of network technologies is **HARD**
  - » Externalities contribute unexpected behaviors
- Core Internet technologies differ from technologies aimed at Internet users
  - » Core Internet technologies: **Scalability and security**
  - » User-facing technologies: **Backward compatibility** and **progressive deployment**
- **Timing matters**: The impact of a late start can be magnified many times (the likely long tail of IPv6 adoption)

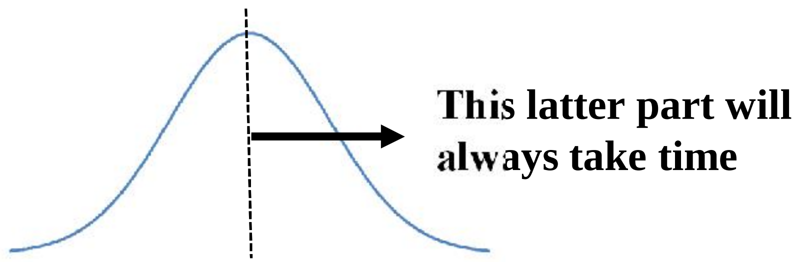

THANK YOU!

References

[1] R. Guerin and K. Hosanagar, "Fostering IPv6 Migration Through Network Quality Differentials." ACM SIGCOMM Computer Communication Review, Vol. 40, No. 3, July 2010.

[2] S. Sen, Y. Jin, R. Guerin, and K. Hosanagar, "Modeling the Dynamics of Network Technology Adoption and the Role of Converters." IEEE/ACM Transactions on Networking Vol. 18, No. 6, December 2010.

[3] M. Nikkhah, R. Guerin, Y. Lee, and R. Woundy, "Assessing IPv6 Through Web Access – A Measurement Study and Its Findings." Proc. ACM CoNEXT 2011 Conference, Tokyo, Japan, December 2011.

[4] M. Nikkhah and R. Guerin, "Migrating the Internet to IPv6: An Exploration of the When and Why." IEEE/ACM Transactions on Networking, Vol. 24, No. 4, August 2016.

[5] M. Nikkhah, A. Mangal, C. Dovrolis, and R. Guerin, "A Statistical Exploration of Protocol Adoption." IEEE/ACM Transactions on Networking, Vol. 25, No. 5, October 2017.